

Dealing with Uncertainty Arising Out of Probabilistic Risk Assessment

Kenneth A. Solomon, William E. Kastenberg,
Pamela F. Nelson

35th
Year

Rand

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE SEP 1983		2. REPORT TYPE		3. DATES COVERED 00-00-1983 to 00-00-1983	
4. TITLE AND SUBTITLE Dealing with Uncertainty Arising Out of Probabilistic Risk Assessment				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Rand Corporation,1776 Main Street,PO Box 2138,Santa Monica,CA,90407-2138				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 64	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

This research was supported by the Oak Ridge National Laboratory under Contract No. 9025.

Library of Congress Cataloging in Publication Data

Solomon, Kenneth A., 1947-
Dealing with uncertainty arising out of probabilistic
risk assessment.

"R-3045-ORNL."
Bibliography: p.
Includes index.
1. Decision-making. 2. Risk. 3. Uncertainty.
I. Kastenberg, W. E. II. Nelson, Pamela F.
III. Rand Corporation. IV. Title.
T57.95.S64 1983 363.1'79 83-17813
ISBN 0-8330-0525-1

The Rand Publications Series: The Report is the principal publication documenting and transmitting Rand's major research findings and final research results. The Rand Note reports other outputs of sponsored research for general distribution. Publications of The Rand Corporation do not necessarily reflect the opinions or policies of the sponsors of Rand research.

R-3045-ORNL

Dealing with Uncertainty Arising Out of Probabilistic Risk Assessment

Kenneth A. Solomon, William E. Kastenberg,
Pamela F. Nelson

September 1983

Prepared for the
Oak Ridge National Laboratory

35th
Year



PREFACE

This report addresses one aspect of how probabilistic risk assessment studies should be integrated into the regulatory process. Intended for use by the Nuclear Regulatory Commission (NRC) staff, it looks specifically at how uncertainty in risk estimates is to be factored into regulatory decisionmaking and is part of a larger ongoing effort. It is funded by the Oak Ridge National Laboratory under Contract No. 9025. A companion study is *An Evaluation of Alternative Safety Criteria for Nuclear Power Plants*, Rand Note N-1806-ORNL, by Kenneth A. Solomon and Pamela F. Nelson, June 1982.

The findings should be of interest to the NRC staff who regulate nuclear energy, to policymakers involved in setting risk acceptance criteria, and to industries that eventually must comply with these criteria.

SUMMARY

The Nuclear Regulatory Commission (NRC), like many other regulatory bodies, often must decide whether some system or some overall technology satisfies a prescribed safety goal. The NRC makes these decisions on the basis of probabilistic risk assessments (PRAs). The objective of a PRA is to quantify the risk of occurrence of some undesirable event, such as a reactor core melt, or of some undesirable consequence of an event, in this example, death. The quantified risk is expressed as a probability, e.g., one core melt expected per 10,000 reactors per year. Estimates of risk are subject to uncertainty. This uncertainty arises from variability in available data on failure rates, difficulties in predicting the effects of unusually stressing events external and internal to the system being assessed, and an insufficient data base on human errors in uncommon situations. Therefore, the actual risk that a part, a system, or a plant will fail may be greater or lesser than the PRA's best estimate of the risk. Indeed, there is associated with each possible risk value a probability that that value is the right one, and it is thus possible to construct a frequency distribution of probabilities over all possible risk values. Past PRAs have equated the best estimate of the risk, that is, the value to be used in comparing the risk to the safety goal, with the median risk value in the frequency distribution, i.e., the median of the set of all possible risk values when weighted by their associated probabilities. It has been suggested that the mean risk value be used instead, because the mean is usually larger than the median in risk-frequency distributions, which generally include some relatively high values with nonnegligible probabilities. Since 50 percent of the frequency distribution of actual risk probabilities falls below the median, systems would have to be designed so that the probability that the actual risk falls below the safety goal is greater than 50 percent if the mean is used as the best estimate.

The NRC has expressed an interest in exploring further the ways in which uncertainty can be taken into account in the regulatory process. In this report, we make several recommendations in that regard, based on

a new approach to the measurement of uncertainty in meeting safety goals and on a review of past decisions made by the NRC and other agencies in setting safety goals.

In our approach we assume that the probability of actual risk of failure is less than the safety goal. We refer to this probability as the "confidence in risk"; it is equal to the fraction of the risk frequency distribution that falls below the safety goal. Thus, if it is required that a system be designed so that the median estimate of risk of failure meets the safety goal, then the confidence in risk is 50 percent. As noted above, if the mean is used instead of the median, we can be more than 50 percent sure that the actual risk falls below the safety goal. However, instead of using the mean, for which the confidence in risk varies with the distribution, we propose that federal regulators consider choosing some level of confidence in risk that satisfactorily accounts for uncertainty. All systems would then have to be designed so that the risk assessor is X percent confident that the risk of failure is less than the safety goal, i.e., that the best estimate of the risk is low enough so that X percent of the risk-frequency distribution falls below the safety goal.

To help us suggest what level of confidence in risk might satisfactorily account for uncertainty, we compared nonnuclear and nuclear safety goals with relevant risk-frequency distributions to infer what levels of confidence had implicitly satisfied regulatory agencies and risk assessors in the past. We constructed the risk-frequency distributions from medians and limits of confidence and from our experience with frequency distributions in analogous circumstances.

We took four approaches to comparing nonnuclear risk-frequency distributions and safety goals. In the first approach, we assumed that median concentrations of carcinogenic pollutants in drinking water supplies and workroom air were equal to standards set by the U.S. Environmental Protection Agency (EPA) and the American Conference of Governmental and Industrial Hygienists (ACGIH). We then inferred median cancer risks from these concentrations, constructed risk-frequency distributions, and compared these with the safety goal for cancer risk set in the NRC's policy statement on safety goals for nuclear power plants (NRC, 1982b). Assuming that environmental carcinogens and

nuclear power plants cause cancer in the exposed population at similar rates, the values of confidence in risk calculated from the above comparison for almost all pollutants imply that the NRC set its safety goal so as to be at least 98 percent sure that the actual cancer risk would fall below it.

In the second approach, we inferred median cancer risks from measured levels of carcinogenic pollutants in drinking water supplies and in the air over Los Angeles, constructed risk-frequency distributions, and compared these with the same NRC safety goal. Assuming, once again, that environmental carcinogens and nuclear power plants cause cancer in the exposed population at similar rates, calculated values of confidence in risk for comparisons involving two-thirds of the pollutants surveyed imply that the NRC set its safety goal so as to be more than 75 percent sure the actual cancer risk would fall below it.

In the third approach, we compared frequency distributions of median cancer risks from measured concentrations of carcinogens in drinking water and workroom air with EPA and ACGIH standards. Values of confidence in risk for comparisons involving two-thirds of the pollutants surveyed imply that safety goals were set so as to be over 90 percent sure that the actual cancer risk would fall below them.

Finally, we compared frequency distributions for failure risk corresponding to yield stresses of structural components with design yield stresses determined by the application of safety factors accepted within the civil engineering profession. The safety factors allow for a confidence in risk of at least 90 percent.

To determine the confidence in risk implicitly deemed acceptable for nuclear power plants by the NRC at present, we used data given in three PRAs to construct frequency distributions for risk of reactor core melt, of acute (or prompt) death in the exposed population, and of latent (cancer) death. We compared these distributions with the NRC safety goals and found that for comparisons involving two of the three PRAs reviewed, we can be over 50 percent confident that the actual risk of core melt falls below the safety goal, at least 97 percent confident that the risk of acute death falls below the goal, and over 99 percent confident in the safety goal for latent death.

On the basis of our review, we offer four recommendations on the treatment of uncertainty in the regulatory process. The first two are procedural in nature and would ensure that sufficient analysis is performed in PRAs to allow implementation of the third recommendation, which suggests a level of confidence in risk that NRC might require, as inferred from the review described above; the fourth recommendation suggests a required level of confidence in risk in instances where the PRA contains insufficient analysis.

Recommendation Number 1: The NRC should require that PRAs specify statistical distributions for each measure of risk (or, at a minimum, upper and lower confidence limits) so that a level of confidence in risk can be determined.

Recommendation Number 2: The NRC should require sensitivity studies to examine the effects of various assumptions regarding certain parameters important to risk and its statistical distribution. These parameters should include contributions from external as well as internal initiators.

Recommendation Number 3: The NRC should specify a minimum level of confidence in risk at least as high as that we have inferred from nonnuclear risks. Such a level might be 90 percent for core-melt frequency and 95 percent to 99 percent for prompt and latent deaths. This high level of confidence would account for unquantifiable risks.

Recommendation Number 4: In the absence of a statistical distribution, the NRC should specify a minimum safety factor, that is, a minimum ratio of safety goal to best-estimate risk. At this point, we lack sufficient information to develop a minimum safety factor for all measures of risk; however, we believe that a factor of 2.5 to 5.0 might be appropriate for core-melt frequency.

In the draft implementation and evaluation plans (contained in NUREG0880), the NRC staff considers the safety goals as design objectives for new plants. For existing plants, less conservative safety goals would be established. In both cases, the staff would require only that the median of the quantifiable risk distribution not exceed the safety goal, i.e., that a 50 percent confidence in risk would be enough. Further risk reduction would be required only if it were cost-effective.

Our procedural recommendations 1 and 2 do not conflict with the staff's plans, although we advocate a more rigorous analysis of risk. However, the high levels of confidence we suggest in recommendations 3 and 4 are more conservative than the Staff is willing to accept at this time.

CONTENTS

PREFACE	iii
SUMMARY	v
FIGURES	xiii
TABLES	xv
Section.	
I. INTRODUCTION	1
Background	1
Scope of the Work	4
II. DEFINITIONS, ASSUMPTIONS, AND PRELIMINARY DISCUSSION	5
Risk and Uncertainty	5
Safety Goals, Risk, and Confidence	7
Inferring a Level of Confidence	8
Relevance to Reactor PRAs	9
Assumptions	9
III. INFERRING CERTAINTY FROM NONNUCLEAR POWER RISKS	11
Introduction	11
Comparison of Chemical Standards and NUREG0880	
(Approach 1)	12
Comparison of Chemical Levels and NUREG0880 (Approach 2) ..	13
Comparison of Chemical Risk Levels and Chemical Standards	
(Approach 3)	16
Comparison of Design Standards and Failure (Approach 4) ..	17
Discussion of Results	20
IV. INFERRING CERTAINTY FROM NUCLEAR POWER RISKS	21
Introduction	21
The NRC Safety Goal	21
Results of the Zion/Indian Point PRAs	23
Estimate of Confidence in Risk	25
Discussion of Results	32

V. FINDINGS AND IMPLEMENTATION GUIDES	35
How To Treat Uncertainty	35
Determining a Reasonable Level of Confidence?	37
The Area Under the Curve Is Not the Only Parameter	38
Ways To Implement Our Findings	40
Appendix	
A. DETERMINATION OF POINT RISK	43
B. DETERMINATION OF CONFIDENCE	44
BIBLIOGRAPHY	49

FIGURES

1.	The Confidence in Risk Estimate Concept: Risk of Core Melt per Reactor Year	7
2.	Confidence Values: Core Melt, Log-Normal Distribution	26
3.	Confidence Values: Acute Deaths, Log-Normal Distribution ...	27
4.	Confidence Values: Latent Deaths, Log-Normal Distribution...	28
5.	Confidence Values: Core Melt, Weibull Distribution	29
6.	Confidence Values: Acute Deaths, Weibull Distribution	30
7.	Confidence Values: Latent Deaths, Weibull Distribution	31

TABLES

1. Summary of Approaches	12
2. Certainty Values Inferred from Approach 1: Chemical Standards Compared with NUREG0880	14
3. Certainty Values Inferred from Approach 2: Chemical Risk Compared with NUREG0880	15
4. Certainty Values Inferred from Approach 3: Chemical Risks Compared with Chemical Standards	17
5. Confidence in Risk Estimate for Nuclear Power Plants	32

I. INTRODUCTION

BACKGROUND

Risk is an important factor in many public policy decisions and regulatory actions. Quantitatively, risk may be defined as the combination of the frequency of occurrence of some undesirable event, such as a nuclear reactor core melt, and its consequences, in this example, death. Risks may be estimated through the techniques of probabilistic risk assessment (PRA). Uncertainty as to the correctness of this estimate can itself be quantified by PRA and can be represented by a probability distribution function over the range of possible risk values.

Numerous government agencies--federal, state, and local--are expanding the role for PRAs in the analysis of technological risks. The Nuclear Regulatory Commission (NRC) has a number of both existing and potential uses for PRA in the regulatory, licensing, and standards development processes (see, for example, Ernst, 1982). The Environmental Protection Agency (EPA) has stimulated the use of PRAs among electric utilities in documenting such things as compliance with the Clean Water Act. The Federal Aviation Administration (FAA) is beginning to use PRA to address such issues as air safety in various postulated air traffic control scenarios. The Food and Drug Administration (FDA) supports the use of PRA in documenting the intent of the Delaney Clause, which states that no food additive should offer any incremental cancer risk above background. This list of examples goes on.

Although we can uncover a vast literature detailing the approaches and methods to be used in estimating both risk (e.g., NRC, 1982a) and uncertainty (e.g., Bolten et al., 1983), we find little written about how regulatory decisions account for uncertainty. Generally, the greater the complexity in the technology under assessment, the greater the uncertainty in the risk estimate. And, as this uncertainty increases, so does the difficulty of the decisionmaking process.

Agencies are just beginning to recognize the need for explicit guidance when acting upon risk assessments where uncertainty plays a major role. In early 1982, the NRC published a proposed policy statement on safety goals for nuclear power plants (NRC, 1982b).¹ Proposed qualitative goals and associated numerical guidelines for nuclear power plant accident risks were presented for public comment along with a supporting discussion paper. As part of the policy paper the Commissioners of the NRC requested that the NRC staff develop a specific action plan for implementing the qualitative safety goals and their associated numerical guidelines. Included was the question: "What further guidance, if any, should be given for decisions under uncertainty?"

The NRC, typical of other regulatory agencies, continually faces the dilemma of making safety related judgments in an arena of varying qualities of information. The implementation of the proposed safety goals and numerical guidelines depends, in a large measure, on PRA. PRA, however, is limited by uncertainties inherent in the underlying data, uncertainties in the underlying physical phenomena, and uncertainties in the methodology itself. Although some of this uncertainty may be reduced by the development of better methodology, the infrequency of nuclear reactor accidents will restrict the growth of the data base, so that overall uncertainties are likely to remain large. As a result, it will not be possible to verify the compliance of nuclear power plants in meeting safety goals because accidents are such low-frequency events.

In its discussion paper, the Commission stated that:

Implementation of the proposed numerical guidelines for reactor safety must take into account this fundamental limitation of PRA [the small data base and consequent inability to verify compliance] and its associated uncertainty.

¹This policy statement is most commonly referred to as NUREG0880, and we will follow this convention in the remainder of our study.

As a first step, the Commission made the following recommendation:

In probabilistic risk assessments made in conjunction with safety goals, the underlying assumptions and associated uncertainties should be disclosed and documented for consideration in the regulatory process. In most situations, these probabilistic risk assessments should be performed during the trial period on the basis of realistic assumptions and best-estimate or mean-value analysis, and they should include an understandable presentation of the magnitude and nature of the uncertainties.

How to factor the magnitude and nature of the uncertainties into regulatory decisionmaking has already become a matter of some controversy. The NRC proposed in NUREG0880 that the median risk in the probability distribution of risk be used to determine whether a system meets the safety goal. Under this proposal, a system would have to be designed so that the probability of failure above the safety goal is no more than 50 percent. In its review of the NRC draft implementation plan, the Advisory Committee on Reactor Safeguards (ACRS) recommended that the mean risk value be compared to the criterion. Since the mean risk estimate is usually higher than the median, using the mean would increase the likelihood that the risk of failure is actually less than the safety goal.

Rather than attempt to resolve the issue of median versus mean as a criterion, our objective is to propose an alternative approach. Given the probability distributions associated with risk estimates, the problem is, then, determining how confident we should be that the goal is met. Or, how to treat the general question of uncertainty in the regulatory process.

To address these questions, we attempt to infer a minimum level of confidence by reviewing a variety of technologies involving risk, their associated uncertainties, and various prescribed safety or acceptance goals. Basically, we identify or deduce statistical distributions for various estimates of risk and then determine the confidence level with which either the goals of NUREG0880 or other goals (e.g., those of the EPA) are met.

SCOPE OF THE WORK

We achieve our fundamental objective of providing guidance on the consideration of uncertainty within the nuclear regulatory process as follows: In Sec. II, we define "confidence in risk"--a simple measure of the degree to which a goal or regulation is likely to be met. We deliver some assertions and assumptions regarding comparative risks in the balance of that section, and in Sec. III, we apply the confidence in risk measure to some selected nonnuclear technologies.

In Sec. IV, we examine some recent PRAs for the Zion and Indian Point nuclear power plants; compare the estimated risks to proposed safety goals; and apply the confidence in risk measure to the frequency of core melt, and the societal risk in terms of acute or prompt deaths and latent (cancer) deaths.

In Sec. V, we develop some decision rules for dealing with uncertainty. These decision rules are based on the results of Secs. III and IV. Finally, we offer some recommendations regarding the implementation of our findings.

II. DEFINITIONS, ASSUMPTIONS, AND PRELIMINARY DISCUSSION

RISK AND UNCERTAINTY

Risk can be expressed either as the probability of occurrence of some event and its consequences, or as the magnitude of the consequences expected over some period of time. Such consequences may include the number of injuries, an amount of money lost, prompt deaths, latent deaths, and property damage, e.g., the amount of land contaminated. Frequency is generally given as the number of undesired events over some period of time, such as a reactor year. For some events, such as those described by actuarial data, the risk may be obtained by simply multiplying the consequences and the frequency. It can be expressed as individual or societal risk. For others, the risk may be estimated and represented by either a mean value or a statistical distribution. These estimates are usually based on models and supported by empirical data.

For nuclear power plants, estimates of risk include various inherent uncertainties. In this work it is convenient to group uncertainties as follows:

Group 1 includes uncertainty as to the representativeness of historical failure rate data for systems and components. Group 1 also includes uncertainty arising from the variability of test data and from the various human errors in operation and maintenance that are reflected in the historical failure rate data. These uncertainties can usually be quantified by a PRA.

Group 2 uncertainties typically arise out of the:

- a. Difficulties in predicting the failure rates of structures (e.g., pressure vessels).
- b. Inability to accurately quantify the effects of severe external events (e.g., seismic shock, depth of floodwater) beyond the "design basis."
- c. Inability to accurately predict (extrapolate) the behavior of systems beyond the conventional "design basis" (e.g., containment leak rates following core melt).

- d. Difficulties in defining the parameters (probability distributions) for certain physical phenomena (e.g., hydrogen detonation, molten fuel retention).
- e. Insufficient quantitative data base of human errors in relatively uncommon situations but which might occur during plant lifetime.

Risks associated with Group 2 uncertainties are usually treated parametrically, i.e., sensitivity studies are performed to determine their influence on total risk.

Group 3 uncertainties are associated with events that could challenge the safety related systems of a nuclear power plant, which, although they cannot be described in detail, have a sufficiently high probability of occurrence to warrant consideration. These uncertainties evolve from:

- a. Design errors.
- b. Unsuspected common-mode failures (e.g., simultaneous failures of redundant safety systems).
- c. Unpredictable human actions, including those of both the operators and saboteurs.

The statistical distributions that represent various measures of risk usually consider uncertainties that are empirically quantifiable or can be estimated parametrically. These include the Group 1 and Group 2 uncertainties above. Group 3 uncertainties are not easily quantified or, as in the case of sabotage, are unquantifiable with present methods. Early PRAs (such as those outlined in NRC, 1975), as well as some recent ones (such as Philadelphia Electric Co., 1981) deal only with Group 1 uncertainties and yield only mean-value risks. Moreover, they ignore such things as external initiators. The Zion and Indian Point PRAs deal with both Group 1 and Group 2 uncertainties. In this work we consider both quantifiable (Group 1) uncertainties and those that can be parameterized and estimated (Group 2), resulting in a statistical

distribution for each measure of risk. From these statistical distributions, a best-estimate or mean-value of the risk can be obtained.

SAFETY GOALS, RISK, AND CONFIDENCE

Graphically we can describe the relationship between mean value (best estimate), degree of confidence, and a safety goal as follows. Consider a hypothetical PRA leading to the statistical distribution for core-melt frequency as shown in Fig. 1. The distribution includes Group 1 and Group 2 uncertainties and ranges from 5.0×10^{-4} to 1.0×10^{-5} core-melt events per reactor year with a best estimate risk of 7.0×10^{-5} events per reactor year. Superposition of a core-melt safety goal of 1.0×10^{-4} per reactor year (from NUREG0880) defines two areas. Area A can be considered a "degree of confidence" with which the risk in question meets the goal. Mathematically, the numerical value of area A is the probability that the goal is met and is equal to the ordinate of




Fig. 1 -- The confidence in risk estimate concept:
risk of core melt per reactor year

a point on a cumulative distribution function derived by integration of the statistical distribution given in Fig. 1. In this work we call A the estimates of confidence in risk. Note that in this example, the best-estimate satisfies the safety goal and, if the distribution is symmetrical, the mean and median are equal. If the NRC were to require that all reactors meet the safety goal with a confidence of X percent, and if A in Fig. 1 were less than X, the reactor design would have to be modified to shift the distribution function far enough to the left to bring A down to X.

We should mention that the value of A depends on two factors: the width of the distribution around the best-estimate value, and the absolute value of the best estimate relative to the standard. Narrow distributions are associated with high-frequency risks because some actuarial data are available. In such cases, the required value of A could be small because even if the mean value is close to the goal, the actual risk is not likely to be much over. In contrast, wide distributions tend to be characteristic of high-consequence, low-frequency events. In these cases, the required confidence A may be great, to ensure that the separation between the best estimate and the goal is large and thus to minimize the probability that the actual risk is much over the goal.

The compound question we then consider is:

1. What values of A are regarded as acceptable for a variety of regulated substances and nonnuclear technologies?
2. And, under what circumstances might these values be applicable to decisions relevant to nuclear reactor PRAs and safety goals?

INFERRING A LEVEL OF CONFIDENCE

We infer values of A (the area under the curve) using four approaches or comparisons. Each comparison is detailed and implemented in Sec. III. Two of them involve latent risk, one involves acute risk, and the other involves property-damage risk. The approaches consist primarily of comparing risks at some estimated or assumed level with appropriate standards for specific chemicals or with standards proposed by the NRC (NUREG0880).

Since we are not interested in high-frequency risks, we eliminate from consideration risks larger than individual mortality rates with best estimates greater than 1.0×10^{-4} per person year (viz., prompt fatalities resulting from automobile accidents, falls, and fires).

RELEVANCE TO REACTOR PRAS

The extent to which we can infer values of A from a variety of regulated substances and various technologies and apply these values to decisions involving uncertainties in nuclear power plant PRAs is limited by our ability to demonstrate similarity between the potential undesirable consequences of these technologies and those of nuclear power. For example, values of A for regulated carcinogenic substances should reflect reasonably well a value for low-level radiation, which is also a carcinogen. Accidents involving property damage, such as a structural failure, might be comparable to core melt.

Values of A for regulating the nuclear industry should be at least as large as those in nonnuclear industries given the intent of the nuclear industry to remain conservative.

ASSUMPTIONS

To use standards for nonnuclear technologies and their attendant risks, the following assumptions are used in Secs. III and IV.

1. For deducing individual risk, the mean human lifetime is 70 years.
2. Occupational standards are 10 times less stringent than general population standards when general population standards are not explicitly stated (Solomon and Abraham, 1980).
3. Animal data on carcinogenic risks can be extrapolated to human effects according to the method discussed in Bolten et al. (1983).
4. The methodology employed in our selected nuclear reactor PRAs is correct.

5. When only upper- and lower-confidence levels are given, either a two-parameter normal, log-normal, or Weibull distribution applies.

III. INFERRING CERTAINTY FROM NONNUCLEAR POWER RISKS

INTRODUCTION

To infer the level of uncertainty tolerated in various circumstances (or conversely, the level of certainty desired) we use four comparisons as discussed below. For each of these comparisons we consider two measures--some risk either estimated empirically or assumed to occur at the regulated standard value with some specific or generic safety standard. We estimate a "statistical distribution" around this risk and then superimpose this distribution on a specific or generic safety standard. The results obtained in this section, together with the findings of Sec. IV, support the recommendations we develop in Sec. V.

The first comparison assumes that the environmental concentration of a given carcinogenic pollutant is equal to the level set by the relevant standard. We calculate the risk of latent death associated with that concentration and, using available confidence limits, construct a probability distribution around this calculated risk. This distribution represents the uncertainty. We then compare this uncertainty distribution to the safety goal set forth in NUREG0880 (i.e., 0.1 percent of the risk from all cancers) to estimate A.

The second and third approaches compare risks associated with actual environmental concentrations of various chemicals with the guidelines of NUREG0880 (for latent effects again) and with the appropriate chemical standards.

The last approach compares failure rate distributions for civil engineering structures with safety factors.

These four comparisons enable us to estimate the "confidence in risk," A. The first three relate to how uncertainty can be factored into NRC's proposed safety goal relating to latent deaths. The design standards relate to the NRC goal for core-melt frequency.

In Approaches 1 and 2, we assume that the standards in NUREG0880 are applicable to nonnuclear risks, i.e., that nuclear power plants and carcinogenic pollutants create similar cancer risks. Also, the data

given below are drawn from studies with varying methodologies and are thus not entirely comparable. However, since we want to infer only a rough estimate or range of reasonable values for A, this is not a critical drawback. These four approaches are summarized in Table 1.

Table 1
SUMMARY OF APPROACHES

Approach No.	Measure of Risk	Standard of Comparison	Relates to Risk of
1	Chemical standard	NUREG0880	Latent death
2	Actual chemical level	NUREG0880	Latent death
3	Actual chemical level	Appropriate standard safety factors	Latent death
4	Failure rates	Appropriate standard safety factors	Core melt

COMPARISON OF CHEMICAL STANDARDS AND NUREG0880 (APPROACH 1)

In this comparison, the risk associated with the standard for a given chemical carcinogen is compared with the numerical guidelines for latent risks proposed in NUREG0880. We translate the chemical standard into a risk value, assign the best-estimate risk to this measure, and then determine a statistical distribution around this best estimate.

The latent cancer risk standard detailed in NUREG0880 specifies an acceptable level of risk at 0.1 percent of the background cancer risk. This NUREG0880 safety goal equates to about 1.9×10^{-6} cancer deaths per year.

While quantifying the NUREG0880 safety goal is quite straightforward, determining the best-estimate risk from a chemical standard is generally not. We determine these best-estimate risks by examining the risks associated with chemicals regulated by two organizations: the EPA and the American Conference of Governmental and Industrial Hygienists (ACGIH).

The EPA's Water Quality Standards

The EPA guidelines designate both mean and upper- and lower-bound standards. For chemicals in drinking water, the EPA guidelines designate the mean lifetime risk at 1.0×10^{-6} (this corresponds to a mean annual risk of about 1.4×10^{-8} , assuming a 70 year lifetime). When we assume a normal distribution with respect to dose and risk and apply a 95th percentile limit to EPA's prescribed upper and lower bounds (the upper and lower bounds are equidistant from the mean), a statistical distribution of risk probabilities can be defined which we then superimpose on the NUREG0880 safety standard.

The ACGIH Standards on Chemicals in the Workroom

The ACGIH specifies allowable workroom concentrations of toxic pollutants in air. We calculate the risks associated with allowable concentrations of several of these pollutants by applying a multihit dose response model (see Rai and Van Ryzin, 1979, or Bolten et al., 1983) and use the virtual safe dose, the best-estimate risk, and the given 97.5 percent confidence limit to estimate the statistical distribution of risk probabilities. We then superimpose this distribution about the NUREG0880 safety goal.

When information on general population standards is not readily available, we convert occupational standards to general population standards by multiplying the occupational standards by 0.1 according to the observations in Solomon and Abraham (1980).

For the majority of selected chemicals shown in Table 2, the value for A is about 0.98 or 0.99; this translates to a high confidence that the risks corresponding to each of the chemical standards satisfies the NUREG0880 goal. This of course assumes that these chemical risks are considered independently.

COMPARISON OF CHEMICAL LEVELS AND NUREG0880 (APPROACH 2)

In the second comparison we consider actual risks associated with actual environmental levels of specific toxic chemicals rather than just standards. We look at several carcinogens whose concentrations exceed the prescribed standard. Specifically, we assess the risk associated

Table 2

CERTAINTY VALUES INFERRED FROM APPROACH 1: CHEMICAL STANDARDS
COMPARED WITH NUREG0880

Inferred Certainty A	Nature of Risk	Best-Estimate (Annual Individual Mortality)	Reference
0.99	Acrylonitrile ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.98	Arsenic ^b	1.3×10^{-7}	<i>Fed. Reg.</i> , 1980
0.99	Benzene ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.99	Benzedine ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.99	Beryllium ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.99	Carbon tetrachloride ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.99	Chloroform ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.99	Chlordane ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.99	DDT ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.60	DDT ^c	1.5×10^{-6}	ACGIH, 1979
0.99	Dichlorobenzedines ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.99	2,4-Dinitrotolune ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.99	1,2-Diphenyl ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.53	Dieldrin ^c	5.7×10^{-6}	<i>Fed. Reg.</i> , 1980
0.99	Halomethanes ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.99	Heptachlor ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.99	Hexachlorobutadiene ^a	1.4×10^{-8}	<i>Fed. Reg.</i> , 1980
0.99	Vinyl chloride ^b	6.0×10^{-5}	<i>Fed. Reg.</i> , 1980

^aWaterborne.

^bThe standard for arsenic is less stringent than for other water-borne chemicals.

^cOccupational exposure.

with drinking water supplies in several U.S. cities and air over Los Angeles and compare these risks with NUREG0880. Our results are reported in Table 3.

We assume that the health risks associated with air pollution as reported by the Air Resources Board (1981) are best-estimate risks, that one standard deviation is one order of magnitude, and that a normal

Table 3

CERTAINTY VALUES INFERRED FROM APPROACH 2: CHEMICAL RISK
COMPARED WITH NUREG0880

Inferred Certainty A	Nature of Risk	Best-Estimate (Annual Individual Mortality)	Reference
0.63	Arsenic ^a	8.0×10^{-7}	ARB, 1981
0.45	Benzene ^a	4.6×10^{-6}	ARB, 1981
0.78	Carbon tetrachloride ^a	4.4×10^{-7}	ARB, 1981
0.46	Chloroform ^a	3.6×10^{-6}	ARB, 1981
0.47	Dialkylnitrosamines ^a	2.9×10^{-6}	ARB, 1981
0.87	Ethylene dichloride ^a	3.3×10^{-7}	ARB, 1981
0.49	Perchloroethylene ^a	2.2×10^{-6}	ARB, 1981
0.99	Trichloroethylene ^a	1.1×10^{-7}	ARB, 1981
0.99	DDT ^b	1.5×10^{-7}	Solomon et al., 1980
0.95	Dieldrin ^b	7.1×10^{-10}	Solomon et al., 1980
0.76	Vinyl chloride ^b	1.8×10^{-9}	Solomon et al., 1980
0.99	Chlordane ^b	7.1×10^{-10}	Solomon et al., 1980
0.99	Lindane ^b	3.0×10^{-10}	Solomon et al., 1980
0.99	PCB ^b	3.8×10^{-9}	Solomon et al., 1980
0.99	Chloroform ^b	2.1×10^{-7}	Solomon et al., 1980
0.99	Carbon tetrachloride ^b	2.3×10^{-10}	Solomon et al., 1980
0.99	Trichloroethelene ^b	2.7×10^{-8}	Solomon et al., 1980

^a Airborne.

^b Waterborne.

distribution applies. We then superimpose the resulting statistical distribution of risk on NUREG0880 and arrive at values for A that are reported in Table 3.

We estimate the risks of selected carcinogens found in drinking water by extrapolating from animal test data and equating to human subjects according to the method of Rai and Van Ryzin (1979). We consider drinking water samples obtained from eleven cities. Our

concentrations are arrived at by averaging the concentration of each of the selected carcinogens over the sampled cities. For each suspected carcinogen, the Safe Drinking Water Committee of the EPA estimates both lifetime risk and an upper 95 percent confidence band at the "low" dose level. We then convert these risks to annual figures and, using a normal distribution, estimate the confidence that they impose less risk than the standard specified in NUREG0880.

Examining Table 3, we observe small values of A for airborne risks (between 0.46 and 0.87) suggesting a low to moderate degree of confidence that the risks associated with air pollution over Los Angeles would meet the NUREG0880 standard. The values for A for waterborne risks are generally high (most at 0.99) suggesting a high degree of confidence that the associated risks would meet the NUREG0880 standard.

In the table, the extent to which the magnitude of A does not simply depend on the mean risk is a measure of the difference in the statistical distributions of risk among the chemicals. Given that each of the entries assumes a normal distribution, the difference in distribution is due to the width of the distribution (as measured, for example, by the standard deviation).

COMPARISON OF CHEMICAL RISK LEVELS AND CHEMICAL STANDARDS (APPROACH 3)

Here, we assess the extent to which the concentration of carcinogenic substances falls below the level of a set of corresponding chemical standards. This approach differs from the second (shown in Table 3), in that here we compare the risk from five pollutants with the corresponding chemical standards rather than the NUREG0880 standard. For risk from radiation, we compare with NUREG0713 (NRC, 1979).

The chemical risks are derived from a study that summarizes the concentration of carcinogens in the drinking water supplies in eleven U.S. cities. The concentrations of the chemicals shown in Table 4 are averages of the concentrations in each of these cities. Confidence limits were specified in Solomon et al. (1982).

This approach (Table 4) demonstrates a reasonably high value of A for all but occupational exposure to vinyl chloride.

Table 4

CERTAINTY VALUES INFERRED FROM APPROACH 3: CHEMICAL RISKS
COMPARED WITH CHEMICAL STANDARDS

Inferred Certainty A	Nature of Risk	Best-Estimate (Annual Individual Mortality)	Reference
0.98	DDT ^a	1.5×10^{-7}	Solomon et al., 1980
0.94	DDT ^b	6.4×10^{-8}	Solomon et al., 1980
0.99	Dieldrin ^a	5.9×10^{-8}	Solomon et al., 1980
0.99	Vinyl chloride ^a	1.8×10^{-9}	Solomon et al., 1980
0.62	Vinyl chloride ^b	1.3×10^{-9}	Solomon et al., 1980
0.98	Radiation ^b	9.0×10^{-5}	NUREG0713

^aWaterborne.

^bOccupational exposure.

COMPARISON OF DESIGN STANDARDS AND FAILURE (APPROACH 4)

In the design and analysis of civil engineering structures, safety factors are employed to allow for the uncertain nature of the loads, structural properties, and models used. Traditionally, safety load factors (numbers greater than one) were applied to the design; load and resistance factors (numbers less than one) were applied to structural properties to insure conservatism. These factors were based on engineering judgments derived from experience and intuition.

More recently, probabilistic distributions have been used to characterize uncertainties in loads and structural properties. As an illustration, we can consider a single structural component acted upon by a random load. The load induces a stress S which has a statistical distribution. The yield stress of the material is also considered a random variable, denoted by R , and a new random variable $F \equiv R - S$ can be defined. When $F \leq 0$, the structural component fails.

If R and S were defined by normal distributions

$$\bar{F} \equiv \bar{R} - \bar{S}$$

and

$$\sigma_F^2 = \sigma_R^2 + \sigma_S^2$$

where the bar denotes expected value (mean or best estimate) and σ is the standard deviation (see Appendix B). Hence the probability of failure is just

$$P_f = P_r(F \leq 0) = \int_{-\infty}^0 p(f) df$$

where $p(f)$ is a normal distribution whose parameters are F and σ_F .

Because of the uncertainties associated with the values of R and S , for every value of F there is a probability $p(f)$ that that value is less than or equal to zero. This probability function has a normal distribution, and the area P_f under the curve for all values of $F \leq 0$ is the probability that the component will fail. This is analogous to area B in Fig. 1. Area A, the confidence in risk, is analogous to $1 - P_f$, which is called the "reliability."

Before discussing values of P_f or $1 - P_f$, it is convenient to introduce three other useful parameters. The central safety factor C_0 is defined by

$$C_0 \equiv \bar{R}/\bar{S}$$

and measures the ratio of mean strength to mean load. The coefficients of variation are defined by

$$\zeta_R \equiv \sigma_R/\bar{R} \quad \text{and} \quad \zeta_S \equiv \sigma_S/\bar{S}$$

and measure the spread in the distributions. Note that as σ approaches zero, the load or resistance becomes deterministic.

A reliability index β can be defined by the number of standard deviations the mean estimates of F are from $F = 0$. The reliability index β is useful in that it accounts for the relative difference between the mean estimate of R and S as well as the narrowness of the distribution. High values of β can thus be attributed to conservatism in design or to good statistical data.

After examining current civil engineering practice, Hart (1982) derived the following commonly accepted values of P_f and β :

strength failures: $P_f = 10^{-4}$, $\beta \approx 3.5$
service ability failures: $P_f = 10^{-2}$, $\beta \approx 2.0$

For a number of different probability distribution functions, P_f 's on the order of 10^{-2} correspond to safety factors (defined as R/S) of 1.4 to 2.2 and P_f 's on the order of 10^{-4} to safety factors of 2.4 to 3.0, and above.

Ellingwood and Galambos (1982) give a range for the reliability index β based on current criteria for the design of reinforced concrete and steel beams. Values of β range between 2.3 and 4.0 for a wide variety of conditions.

For some load examples, in geotechnical engineering, Meyerhoff (1982) lists the following minimum safety factors:

<u>Item</u>	<u>Safety Factor</u>
Dead loads	0.9 - 1.2
Live loads	1.0 - 1.5
Static water pressure	1.0 - 1.2
Environmental loads	1.2 - 1.4

These values correspond to a 90 percent reliability; i.e., to have a confidence in risk of 0.9, the best-estimate load should be less than its limiting value by 25 percent and 33 percent, for dead loads and live loads, respectively. For stability of earthen structures and foundations, safety factors of 1.9 to 3.3 were found to yield a reliability of 99 percent. Here the goal is set at factors of 2 to 3 above the median value to achieve a confidence in risk estimate of 0.99. Or, to put it another way, for a confidence in risk of 0.90, the load goal and the mean estimate of yield stress should be at least 50 percent apart. For a confidence in risk of 0.99, they should be at least a factor of 3 apart.

DISCUSSION OF RESULTS

We have compared the statistical distribution of some selected nonnuclear risks to various proposed and accepted regulatory standards. Although these statistical distributions are only estimates, they indicate a high (albeit variable) degree of confidence that a specific goal or regulation would be met. When comparing a range of chemical carcinogens and their associated risk distributions with various standards, values in excess of 95 to 98 percent were common for latent deaths.

Safety factors (the ratios of goals to estimated mean risks) for civil engineering structures on the order of 1.0 to 1.2 were associated with a level of confidence of 90 percent. Factors on the order of 2 to 3 were associated with a confidence level of 98 to 99 percent.

In the next section, we estimate the confidence in risk factor A for some nuclear power reactor risk measures. The results will be compared with those obtained here and will be used to formulate our ultimate recommendations.

IV. INFERRING CERTAINTY FROM NUCLEAR POWER RISKS

INTRODUCTION

The assessment of nuclear reactor accident risks is carried out within the framework of probabilistic risk assessment. PRAs have been paying progressively more attention to uncertainty. Although the WASH1400 PRA (NRC, 1975) attempted to estimate the accident risk at nuclear power plants by calculating the frequencies of various accident sequences for two typical reactors and their attendant consequences, the treatment of uncertainty was not emphasized. In the Indian Point (Power Authority, State of New York, 1982) and Zion (Commonwealth Edison Co., 1982) PRAs, an attempt is made to quantify uncertainty. Results are presented for various measures of risk (including core-melt frequency) and uncertainty given in terms of upper (90 percent) and lower (10 percent) confidence bounds. This uncertainty analysis includes both internal and external causes of failure but not sabotage. Because the results of these PRAs are well documented, they are used extensively in this section for illustrative purposes. We emphasize that we made no attempt to evaluate the methodology used in obtaining these results, their confidence limits, or their completeness.

First we briefly summarize the quantitative guidelines contained in the NRC's draft policy statement and their application to the Zion and Indian Point sites. We compare the relevant results of the Zion/Indian Point PRAs with the proposed safety goal. We then use this comparison to generate an "estimate of confidence level" for the frequency of core melt, societal risk in terms of prompt death, and societal risk in terms of latent death. We conclude with a discussion that sets the stage for the decision rules that follow in Sec. V.

THE NRC SAFETY GOAL

In NUREG0880, the NRC has proposed four numerical guidelines:

1. A criterion for the frequency of core melt.
2. Limits on individual and societal risk of prompt death.

3. Limits on individual and societal risk of latent death.
4. A cost-effectiveness criterion in terms of cost allowed per person-rem to decrease exposure to levels beyond compliance with the goals above.

In the development of our decision rules, we employ the core-melt criterion, and the limits on societal risk for both prompt and latent deaths. The cost-effectiveness criterion will be treated in another study. We choose societal risk over individual risk because the available PRAs specify risks for the population in the 50-mile radius surrounding the site. (The methodology developed in this report would apply equally to individual risk.)

The safety goals used in this analysis are:

1. The frequency of core melt should not exceed 1.0×10^{-4} per year.
2. The risk to the population in the vicinity of a nuclear power plant site of prompt fatalities that might result from reactor accidents should not exceed 0.1 percent of the sum of prompt fatality risks resulting from other accidents to which members of the U.S. population are generally exposed.
3. The risk to the population in the area near a nuclear power plant site of cancer fatalities that might result from reactor accidents should not exceed 0.1 percent of the sum of cancer fatality risks resulting from all other causes.

In applying this guideline, the NRC proposes that for societal risk of prompt death, the area within one mile of the plant site boundary be used. For latent deaths, a 50 mile radius should be used. For the Zion and Indian Point sites, these distances can be used to determine the acute and latent risk limits as follows. The individual risk of prompt fatality in the United States, regardless of cause, is about 5.0×10^{-4} per year. Hence the goal requires that the societal risk be less than

$0.001 \times 5.0 \times 10^{-4}$ per year \times population within one mile of the site.

Also the population within one mile of the 111 U.S. nuclear power plant sites ranges between 0 and 1,400 persons, with 168 as an average. Using 168, 500, and 1,000 gives the limit for acute risk as:

$$\begin{array}{lll} 168 & \sim & 0.8 \times 10^{-4} \text{ deaths per year} \\ 500 & \sim & 2.5 \times 10^{-4} \text{ deaths per year} \\ 1,000 & \sim & 5.0 \times 10^{-4} \text{ deaths per year} \end{array}$$

Since the Zion and Indian Point sites are in densely populated areas of the country, we chose 5.0×10^{-4} deaths per year as the limit for prompt death.

Roughly 19 persons per 10,000 U.S. population die of cancer each year. Hence the goal requires that the societal risk of latent death be less than:

$$0.001 \times 19 \times 10^{-4} \text{ per year} \times \text{population within 50 miles of the site}$$

The population within 50 miles of a nuclear plant ranges between 7,700 and 17.5 million. Because the Zion and Indian Point sites are the most populated in the country, we use a population of 17.5 million, yielding a risk limit of 32 deaths per year.

We can summarize these numerical guidelines as follows for the Zion and Indian Point sites:

$$\begin{array}{ll} \text{Limit on core-melt frequency:} & 1.0 \times 10^{-4} \text{ events per year} \\ \text{Limit on societal risk} & \\ \quad \text{of prompt death:} & 5.0 \times 10^{-4} \text{ deaths per year} \\ \text{Limit on the societal risk} & \\ \quad \text{of latent death:} & 32 \text{ deaths per year.} \end{array}$$

RESULTS OF THE ZION/INDIAN POINT PRAS

To derive levels of confidence in risk for nuclear power plants, estimates for core-melt frequency and the societal risk of prompt and latent deaths are derived from the PRAs for Zion and Indian Point.

For Indian Point Units 2 and 3, the median, mean, and upper 90 percent confidence values of core-melt frequency are given in the PRA. For the Zion plant, only the mean value is given for total core-melt frequency. We estimate the median and upper 90 percent confidence limit from the probability distribution function given for externally induced core-melt frequency. For the three plants we have:

Frequency of Core Melt (per year)	Risk of Latent Death (deaths/year)		
	Indian Point 2	Indian Point 3	Zion
Median	4.0×10^{-4}	9.0×10^{-5}	4.0×10^{-5}
Mean (best estimate)	4.7×10^{-4}	1.9×10^{-4}	6.7×10^{-5}
Upper 90 percent	1.0×10^{-3}	5.5×10^{-4}	2.4×10^{-4}

For the Zion and Indian Point sites, the societal risk for five damage indexes¹ are given in the PRAs as a set of complementary cumulative distribution functions (CCDF). Moreover, for each damage index, the value of the CCDF is reported at the 10th, median (50th), and 90th percentile confidence limits. As shown in Appendix A, the point risk (expected value) can be obtained by calculating the area under the CCDF curve. Using a simple numerical procedure (trapezoidal rule) we obtain:

Percentile	Risk of Acute Death (deaths/year)		
	Indian Point 2	Indian Point 3	Zion
0.1	5×10^{-7}	5×10^{-8}	--
0.5	2×10^{-6}	7×10^{-6}	5×10^{-7}
0.9	5×10^{-4}	4×10^{-5}	5×10^{-5}

¹The five are acute fatalities, injuries, thyroid cancer, latent fatalities, and population dose.

Risk of Latent Death (deaths/year)			
Percentile	Indian Point 2	Indian Point 3	Zion
0.1	3×10^{-3}	1×10^{-3}	1×10^{-3}
0.5	1.5	8×10^{-2}	5×10^{-2}
0.9	10	1	5×10^{-1}

We do not evaluate either the methodology or its application in arriving at the risk figures above. We are aware that these values are under review by the NRC staff and are subject to change. Potential changes are not expected to alter our conclusions dramatically.

ESTIMATE OF CONFIDENCE IN RISK

To determine the confidence in risk, we specify the various measures (core-melt frequency, risk) probabilistically. Since statistical distributions are not given, we assume that the various risk measures are log-normally distributed. Moreover, to test the sensitivity to an assumed distribution, we also use the Weibull distribution.

It should be noted that the log-normal and Weibull distributions used are two-parameter distributions and only two data points are necessary for curve fitting. For prompt and latent deaths, three data points are given. Since we are interested in the high confidence end, the median (50 percent) and 90 percent values are used. If the distributions were truly log-normal or (Weibull), the 10 percent values would be close to the curve. Since this closeness occurs only in one case, we know the true distribution is not log-normally distributed.

As shown in Appendix B, the log-normal distribution function has the useful property that when the cumulative distribution function is plotted against the logarithm of the argument on "normal curve" graph paper, a straight line results.

Similarly, if a variable conforms to a Weibull distribution, it becomes a straight line when plotted on "extreme value" probability

paper. As shown in Appendix B, the coordinates are $\log x$ and $\log \{ \ln[1/F(x)] \}$ where $F(x)$ is the complementary cumulative distribution function.

In Figs. 2 to 4 we have plotted the results of the Zion and Indian Point PRAs assuming log-normal distributions, and in Figs. 5 to 7, we have plotted them assuming a Weibull distribution. The confidence in risk estimate is the ordinate of the intersection of the (complementary) cumulative distribution function with the NRC numerical guidelines (safety goal) calculated for this study. Again, the confidence in risk estimate, A , as defined in Sec. II, is the probability that the risk is actually less than the safety goal.




Fig. 2 -- Confidence values: core melt, log-normal distribution




Fig. 3 -- Confidence values: acute deaths, log-normal distribution




Fig. 4 -- Confidence values: latent deaths, log-normal distribution




Fig. 5 -- Confidence values: core melt, Weibull distribution




Fig. 6 -- Confidence values: acute deaths, Weibull distribution




Fig. 7 -- Confidence values: latent deaths, Weibull distribution

Table 5

CONFIDENCE IN RISK ESTIMATE FOR NUCLEAR POWER PLANTS

Reactor	Median Risk	Log-Normal	Weibull
Core-Melt Frequency (Safety Goal: 1×10^{-4} /yr)			
Indian Point 2	4×10^{-4} /yr	.03	.10
Indian Point 3	9×10^{-5} /yr	.52	.53
Zion	4×10^{-5} /yr	.72	.75
Societal Risk of Early Deaths (Safety Goal: 5×10^{-4} d/yr)			
Indian Point 2	2×10^{-6} d/yr	.900	.9000
Indian Point 3	7×10^{-6} d/yr	.999	.9999
Zion	5×10^{-7} d/yr	.970	.970
Societal Risk of Latent Deaths (Safety Goal: 32 d/yr)			
Indian Point 2	1.5 d/yr	.980	.98
Indian Point 3	0.08 d/yr	.999	.9999
Zion	0.05 d/yr	.9998	.99999

DISCUSSION OF RESULTS

Assuming that the risk estimates are accurately described by the log-normal and Weibull distributions, the degree of confidence that the core-melt risk would meet the safety goal is small for Indian Point 2 (3 percent to 10 percent), because neither the mean nor median value meets the goal. For Indian Point 3, the confidence level is approximately 52 percent, reflecting the proximity of the goal to the median and mean (0.9×10^{-4} and 1.9×10^{-4} , respectively). For Zion, the confidence in meeting the core-melt frequency goal level is between 72 percent and 75 percent depending upon the distribution.

For the societal risk goals, we can be 90 percent confident that Indian Point 2 would meet the safety goal for prompt deaths; for Indian Point 3 and Zion we can be 97 percent confident or better. For latent

death, we can be 98 percent confident or better that the safety goal would be met by all three plants.

The above results demonstrate a principle of potential significance in incorporating the confidence in risk approach into the regulatory process: The greater the confidence level, the less it matters which distribution is chosen as an approximation, and for confidence levels above 50 percent, it matters very little.

At this point it is of interest to compare the results of Sec. III with the results presented herein. For the societal risk measures (both acute and latent), the high degree of confidence compares favorably with the risks from a range of toxic chemical species. For core-melt frequency the results are less comparable. In addition to examining the confidence in risk estimate, it is useful to examine other aspects of the statistical distribution of core melt.

The slope of the lines drawn in Figs. 2 through 7 give some further indication of the distribution's width about its median value. The steeper the curve, the closer the upper and lower confidence limits are to the median. Another indication of the width of the distribution is the ratio of the upper 90 percent confidence bound to the mean. In addition, the ratio of the upper 90 percent confidence bound to the goal measures the distance of the distribution from the goal. For the estimates of core-melt frequency we have:

<u>Site</u>	<u>90% Value/Mean</u>	<u>90% Value/Goal</u>
Indian Point 2	2.1	10.0
Indian Point 3	2.9	5.5
Zion	2.9	2.4

For civil structures, the ratio of 90 percent value to mean runs from about 1.2 to 1.5 in most cases.

As an alternative, we can apply the safety factor concept. If we divide the safety goal for core-melt frequency (1.0×10^{-4} per reactor year) by the mean value, an analogous "safety factor" can be defined. These can be compared to the results of Sec. III:

Confidence	Type	Safety Factor (Goal/Mean)
0.03 - 0.10	Indian Point2 ^a	0.2
0.52 - 0.53	Indian Point3 ^a	0.5
0.72 - 0.75	Zion ^a	1.5
0.90	Soil load	1.0 - 1.9 ^b
0.99	Soil stability	1.9 - 3.3 ^b

^aRatio of mean core-melt frequency to safety goal.

^bMean/goal; because the relationship between these two values for nuclear reactors is the inverse of that for civil structures, the inverse safety factors are analogous.

By extrapolating the relationship between confidence and safety factor for nuclear power plants, it appears that a safety factor of ~2.5 corresponds to a confidence in risk estimate of 0.9 for core-melt frequency. Hence, in the absence of a distribution function, we may employ a safety factor to ensure some degree of confidence that the goal is actually met and thus account for uncertainty.

In the next section, we discuss the degree of confidence (confidence in risk estimate) for these nuclear power plants, the slope of the curves, and the median risks and compare them with the risks of nonnuclear technologies as described in Sec. III. On the basis of these two sets of confidence values we develop some decision rules.

V. FINDINGS AND IMPLEMENTATION GUIDES

HOW TO TREAT UNCERTAINTY

NRC's proposed policy statement on safety goals for nuclear power plants contains both qualitative goals and numerical guidelines. The numerical guidelines are singular values; limits on individual and societal risk of prompt and latent deaths, limits on core-melt frequency, and a limit on cost benefit. In its original action plan for implementing the Commission's proposed safety goals,¹ the NRC staff proposed that PRAs should be performed using realistic assumptions, and the estimates of risk made on median values after formulating uncertainty distributions. In its review of the staff's draft implementation plan, the Advisory Committee on Reactor Safeguards (ACRS) recommended that the mean value be used because it is usually larger than the median (for the cases of interest), and this would ensure with a greater degree of confidence that the numerical guidelines were met.

In Sec. IV we presented the median, mean, and upper 90 percent confidence limits for the frequency of core melt calculated for three reactors. For Indian Point 2, neither the median nor the mean met the numerical guidelines. For Indian Point 3 the median satisfied this guideline, but the mean did not. Last, for Zion both the mean and the median met this guideline, but the upper 90 percent value did not.

Rather than trying to resolve the issue of median versus mean, with respect to this guideline, a more appropriate question might be: How certain should we be that this guideline is met?

Of the four recommendations we present, the first two offer the prerequisites to establishing a quantitative response. These first two recommendations are procedural in effect. The third recommendation establishes a quantitative measure of the confidence required in the presence of sufficient PRA data, and the fourth recommendation establishes a quantitative measure in the absence of sufficient PRA data.

¹On January 7, 1983, the staff issued a revised document, "Plan for Evaluating the Proposed Safety Goals," which has a different focus as implied by the title.

Recommendation 1: The NRC should require that PRAs specify statistical distributions for each measure of risk (or, at a minimum, upper and lower confidence limits) and that their distributions be used to determine a level of confidence in risk.

At the present time, many PRAs fail to detail uncertainty in each of the risk measures. If a quantitative level of confidence is to be determined, uncertainties that are quantifiable should be propagated through PRAs and the risk should be presented as statistical distributions rather than as median values. As stated in Sec. II, some contributions to risk have large uncertainty. Since the results of these PRAs will play a key role in value-impact assessment (the cost-benefit analysis of risk reduction alternatives) they will be an important input into regulatory decisionmaking. If only median values are presented, our knowledge of risk, and hence value-impact, will be masked, and of little use. The NRC staff's original implementation plan, as well as its new evaluation plan, advocate consideration of uncertainty, but conformance with safety goals is based on median values without regard to confidence.

Uncertainties that are not easily quantified, such as those arising from the response to severe external events (beyond the design basis), the behavior of systems in severe accidents (beyond the design basis), and post core-melt fission product behavior and human errors, are usually treated parametrically. To determine the influence of these uncertainties on various measures of risk, and hence on value-impact considerations, we make the following recommendation:

Recommendation 2: The NRC should require sensitivity studies to examine the effects of various assumptions regarding certain parameters important to risk and its statistical distribution. These parameters should include contributions from external as well as internal initiators.

Recommendation 2 is only partially consistent with the NRC staff's plans that call for appropriate sensitivity analyses for certain parameters important to risk. The staff's approach is based on the assumption that the contribution to risk from external events is normally small, provided that the NRC's deterministic requirements are met. Hence, in the staff's view, these contributions can be neglected and as a result, the major portion of the goal can be allocated to internally induced transients and loss-of-coolant accidents.

Recent PRAs, such as those for Zion and Indian Point, show that although the contribution of external initiators to core-melt frequency is smaller than internal initiators, their contribution to public health risks is dominant. This occurs because the consequences associated with seismically induced core melts, fires, and other external events are very large and outweigh their lower frequency. It is in this light that we make Recommendation 2.

DETERMINING A REASONABLE LEVEL OF CONFIDENCE?

Recommendation 1 raises the question: "What level of confidence is reasonable?"

In Sec. IV, we saw that given a statistical distribution, or upper and lower confidence levels (so that a distribution might be derived), a level of confidence in meeting a safety goal could be obtained. For Indian Point 2, for example, we saw that even though the median and mean calculated core-melt frequencies were above the safety goal, there was still a 3 percent to 10 percent chance the goal was met. For Zion, on the other hand, both the median and the mean were below (safer than) the goal, but there was still a 25 percent chance the goal was not met.

For the societal risk of prompt fatalities, the level of confidence in meeting the goal varied between 90 percent and 99.99 percent. For latent fatalities, the level of confidence in meeting the goal varied between 98 percent and 99.999 percent.

In an effort to determine what constitutes a reasonable level of confidence, we examined confidence limits for nonnuclear risks. As discussed in Sec. III, we used a confidence in risk estimate for some comparable risks and displayed them in Tables 1 through 5, inclusive. On the basis of Secs. III and IV, we next recommend:

Recommendation 3: The NRC should specify a minimum level of confidence in risk at least as high as that we have inferred from nonnuclear risks. Such a level might be 90 percent for core-melt frequency and 95 percent to 99 percent for prompt and latent deaths.

This high level of confidence would reflect conservatism to account for uncertainties that are not easily quantified or may be unquantifiable. As such, Recommendation 3 is a departure from the NRC staff's draft implementation plan and draft evaluation plan, which simply call for the median value to be used. The staff does recognize that for some accident sequences, an upper-bound estimate might be so large that regulatory action is appropriate, but the staff would condition such actions on a favorable cost-benefit analysis.

THE AREA UNDER THE CURVE IS NOT THE ONLY PARAMETER

While the value of A might be a good indicator of the relative certainty achieved by varying risks with a fixed distribution and a fixed standard, different distributions might achieve identical values of A yet convey a very different message regarding uncertainty. For example, we can conceive of two distinct distributions, one very flat and the other very peaked, both demonstrating the same value of A yet one conveying more precision. Or, we may have a situation where the value of A may be less, but the statistical distribution is peaked, yielding a "narrower confidence band." In this case, the regulator might tolerate a smaller value of A, if it could be demonstrated that the risk at the 95 percent to 98 percent confidence limits were very close to the median (50 percent) risks.

We may wish to consider other parameters, such as a reliability index or a safety factor, which are related to the width of the statistical risk distribution. In Sec. IV, we saw that the slope of the (complementary) cumulative distribution function gave a measure of the narrowness of the distribution. We can identify two cases of interest. In the first, we have a high confidence (say, 95 percent to 98 percent) that the goal is met, but the slope is such that there is a large

increment in risk for a small increment in confidence. In the second, the confidence is smaller that the goal is met (the median nearly equals the goal) but there is a small increment in risk for a large increment in confidence. The NRC might prefer the first case and condition its confidence requirement on the slope of the cumulative distribution function.

We have searched the nonnuclear data presented in Sec. III and can provide only a limited basis for making a quantitative recommendation as to an alternative criterion to account for uncertainty. The reliability index discussed in Sec. III is a useful quantity because it includes both the relative width of the distribution and the separation between the goal and the mean. The reliability index would have been less than one for all three nuclear cases examined, and less than zero for two of them, whereas current civil engineering practice yields values on the order of 2 to 4 for various structures. The safety factor discussed in Approach 4, however, gives a relationship only between a goal and the mean value independent of the spread or variance. As such, we would expect, for safety situations whose variance in outcome is similar to structural failures following soil load and soil stability mishaps, that a safety factor between 1.0 and 3.0 might be reasonable. However, because the variance associated with the frequency of a core melt is higher,² we would want to require a larger safety factor for core melt. While we cannot yet select a precise value for this safety factor, values between 2.5 and 5 appear to be comparable to 90 percent to 99 percent confidence. We are then led to our next recommendation.

Recommendation 4: In the absence of a statistical distribution, the NRC should specify a minimum safety factor (that is, a minimum ratio of safety goal to best-estimate or mean risk). At this point, we lack sufficient information to develop a minimum safety factor for all measures of risk; however, we believe that a safety factor of 2.5 to 5.0 might be appropriate for core-melt frequency.

²That is to say, the uncertainty distribution for the core-melt situation is likely to be wider (flatter) than the corresponding distribution in the soil load and soil distribution cases.

Recommendation 4 is intended to be a mechanism for providing a satisfactory level of confidence in risk in the absence of a statistical distribution. This approach is somewhat more conservative than that advocated in the staff's draft implementation plan.

WAYS TO IMPLEMENT OUR FINDINGS

The NRC staff's original implementation plan considers the numerical guidelines in NUREG0880 to be design objectives that must be achieved in the design of new plants. For existing plants, less conservative operational levels would be established and cost-benefit considerations would be used to determine whether or not the design objective should be met. The median risk would be compared with either the design or operational objective after propagation of quantifiable uncertainties. Regulatory action would be determined depending upon the type of plant and its status. For example, a plant under construction would have to meet all operating limits using median estimates from a PRA, without regard to cost-benefit guidelines, and should meet all design objectives subject to cost-benefit.

To understand how to implement our recommendations, it is useful to compare them with the staff's implementation plan summarized above. Recommendations 1 (the need for distributions and the specification of uncertainty limits) and 2 (the need for sensitivity studies) involve procedures and are consistent with the staff's approach. However, we advocate broader, more inclusive implementation as follows. While safety goal PRAs may not be required for operating reactors and those awaiting operating license review, a number of PRAs have been completed (Zion, Indian Point, and Limerick) because they are located at high population density centers, or because they were part of the Interim Reliability Evaluation Program. These PRA studies are important because they represent the current status of knowledge concerning the risks of plants licensed (or about to be licensed) in the United States. As such, they can provide information on the importance and contribution of different types of uncertainty, and on those areas of plant design and operation most amenable to possible improvement. Furthermore, they can be used to assess the need for regulatory action based on value-impact consideration.

For the PRAs mentioned above, existing methods can be used to quantify uncertainty resulting from data variability and related factors. As part of the staff's evaluation plan, the suggested single monetary value of averted person-rem for implementing the benefit-cost guideline will be examined, as well as other measures. The identification of proper sensitivity studies is a difficult task, requiring at a minimum an empirical basis.

As part of the construction permit review, new plants will require a safety goal PRA. Pending the results of the staff's evaluation, we advocate full implementation of Recommendations 1 and 2 for these safety goal PRAs.

Recommendations 3 (confidence levels) and 4 (safety factors) involve judgment and can be implemented as follows. For existing PRAs, it may be difficult to provide statistical distributions for all measures of risk (e.g., for both internal and external contributors), and to perform adequate sensitivity studies. To determine whether or not regulatory action is required for these cases, a safety factor could be defined and used to allocate risk. For example, the Limerick PRA estimated a core-melt frequency of 1.5×10^{-5} per reactor year resulting from internal events only. If this is taken as a mean value, the resultant safety factor of 6.66 is large enough to account for uncertainties arising from quantified internal events at the 99 percent level and still leave room to account for uncertainties from external events too. Brookhaven National Laboratory, in its review of the Limerick PRA, gives a value of 1.1×10^{-4} per reactor year for the mean value. If this is the case, there is no margin for uncertainty (the safety factor is less than one) and regulatory action, with or without cost-benefit, might be appropriate.

As stated above, safety goal PRAs will be required for new plants. For these plants, the design objective for the various safety goals should be met at a high confidence level. If this is the case, no further reduction in core-melt frequency would be necessary. This contrasts to the staff viewpoint, wherein the design objectives are met at the median value, and further reductions are subject to cost-benefit guidelines.

In summary, implementation of the recommendations put forward in this work is based upon a high level of confidence rather than a median value. If these levels are met, no further risk reduction would be required. In the staff's view, meeting the goal at the design level would not preclude cost-benefit considerations of further risk reduction.

Appendix A
DETERMINATION OF POINT RISK

The complementary cumulative distribution function $F(x \geq X)$, which represents the results of various probabilistic risk assessments, can be viewed as the frequency of events with consequences x that are greater than a given number X . The definition of $F(x \geq X)$ is just

$$F(x \geq X) = \int_X^{\infty} f(x) \, dx = 1 - \int_{-\infty}^X f(x) \, dx \quad (\text{A.1})$$

where $f(x)$ is the probability distribution function for x . The expression $f(x) \, dx$ is taken as the frequency with which events with consequences x and $x + dx$ occur. The risk is defined as

$$R = \int_{-\infty}^{+\infty} x f(x) \, dx \quad (\text{A.2})$$

Since $f(x) = -\partial F / \partial x$, we can integrate Eq. (A.2) to yield

$$R = \int_{-\infty}^{+\infty} F(x \geq X) \, dx \quad (\text{A.3})$$

Hence the expected value or risk is the area under the curve given by the complementary cumulative distribution function.

Appendix B
DETERMINATION OF CONFIDENCE

INTRODUCTION

To determine the confidence with which a safety goal or criterion is met, the probability distribution function associated with the regulated risk must be known. Usually, only the mean value is given, or attainable. For some specific cases, mean, median, and upper and lower confidence values may be given. While there is no a priori reason to expect that either measured or calculated risks would obey a specific probability distribution function, we assume for the sake of illustration that they do. For this work the log-normal and Weibull distributions are used. In this appendix, some properties of each are discussed.

LOG-NORMAL DISTRIBUTION

A random variable x has a log-normal distribution if $y = \ln(x)$ is normally distributed, i.e.,

$$f(y) = \frac{1}{\sqrt{2\pi}\sigma} \exp \left[-\frac{1}{2} \left(\frac{y - m}{\sigma} \right)^2 \right] \quad (B.1)$$

$-\infty \leq y \leq \infty$

where m = mean of y and σ^2 the variance. Here σ is called the standard deviation. Using a change of variable [$y = \ln(x)$] and noting that

$$p(x) dx = f(y) dy \quad (B.2)$$

we obtain the log-normal probability distribution function $p(x)$ defined by

$$p(x) = \frac{1}{x\sqrt{2\pi}\sigma} \exp \left[-\frac{1}{2} \left(\frac{\ln(x) - m}{\sigma} \right)^2 \right] \quad (B.3)$$

By definition, the expected value and variance of $\ln(x)$ are:

$$E[\ln(x)] = m, \quad (B.4)$$

$$\text{Var} [\ln(x)] = \sigma^2 \quad (B.5)$$

It can be shown that the expected value and variance of x are:

$$E(x) = E(e^y) = e^{m+\sigma^2/2} \quad (B.6)$$

$$\text{Var} (x) = \text{Var} (e^y) = e^{(2m+\sigma^2)} (e^{\sigma^2} - 1) \quad (B.7)$$

The probability that a log-normally distributed variable *does not* exceed a given value X is given by the cumulative distribution function

$$F(X) = \text{Pr} (x < X) = \phi\left(\frac{\ln(X) - m}{\sigma}\right) \quad (B.8)$$

where

$$\phi(\eta) = \frac{1}{(2\pi)^{1/2}} \int_{-\infty}^{\eta} e^{-t^2/2} dt \quad (B.9)$$

Note that when $\ln(X) = m$, the argument of ϕ is zero. Furthermore, $\phi(0) = 1/2$, since m is the mean value of y , and $y = \ln(x)$. Thus e^m (the geometric mean) is the median value of a log-normally distributed variable.

The log-normal distribution has the useful property that when the distribution function $F(X)$ is plotted against the logarithm of its argument on "normal curve" graph paper, a straight line results. As shown above, $F(X) = 0.50$ at the median value m_g . Furthermore, when $F(X) = 0.84$, $\ln(X) = \ln(m_g) + \ln(\sigma_g)$ where

$$\sigma_g = e^{\sigma} \quad (B.10)$$

is the geometric standard deviation. Similarly, $F(X) = 0.16$ when $\ln(x) = \ln(m_g) - \ln(\sigma_g)$ and so on.

WEIBULL DISTRIBUTION

A random variable x has a Weibull distribution if

$$f(y) = \left(\frac{\lambda}{\beta}\right)\left(\frac{x}{\beta}\right)^{\lambda-1} \times \exp \left[-\left(\frac{x}{\beta}\right)^{\lambda}\right] \quad (\text{B.11})$$

where λ and σ are two parameters related to the mean and variance by

$$m = \beta \Gamma(1 + 1/\lambda) \quad (\text{B.12})$$

$$\sigma = \beta^2 [\Gamma(1 + 2/\lambda) - \Gamma^2(1 + 1/\lambda)] \quad (\text{B.13})$$

The probability that a variable governed by the Weibull distribution exceeds a given value X is defined by the complementary cumulative distribution function

$$\bar{F}(X) = \exp \left[-(X/\beta)^{\lambda}\right] \quad (\text{B.14})$$

where

$$\Pr (x \geq X) = \bar{F}(X) \quad (\text{B.15})$$

Taking the logarithm of Eq. (B.14) and changing sign yields

$$\ln[1/\bar{F}(x)] = (X/\beta)^{\lambda} \quad (\text{B.16})$$

Taking the logarithm of both sides yields

$$\log \{\ln[1/\bar{F}(X)]\} = \lambda(\log X - \log \beta) \quad (\text{B.17})$$

Equation (B.17) will be a straight line on "extreme value" probability paper where the coordinates are $\log x$ and $\log \{\ln[1/\bar{F}(x)]\}$.

Note that λ is the slope of the line. When $X = \beta$,

$$\log [\ln(1/\bar{F})] = 0 \quad (\text{B.18})$$

or

$$F = e^{-1} = 0.368$$

BIBLIOGRAPHY

- ACGIH (American Conference on Governmental and Industrial Hygienists), *Threshold Values for Chemical Substances in Workroom Air*, Washington, D.C., 1979.
- Albert, Roy E, "The Carcinogen Assessment Group's Final Risk Assessment on Arsenic," unpublished document prepared at the request of the EPA Office of Air Quality Planning and Standards, Washington, D.C., May 1980.
- BEIR (Biological Effects of Ionizing Radiation), *The Effects on Populations of Exposure to Low Levels of Ionizing Radiation*, National Academy Press, Washington, D.C., 1980.
- Bolten, J. G., P. F. Morrison, and K. A. Solomon, *Risk-Cost Assessment Methodology for Toxic Pollutants from Fossil Fuel Power Plants*, The Rand Corporation, R-2993-EPRI, June 1983.
- Brookhaven National Laboratory, "A Review of the Limerick Generating Station PRA," Draft Report, New York, 1982.
- California Air Resources Board, *Air Resources Board Pollution Study of 1981*, Sacramento, June 1981.
- Commonwealth Edison Co., *The Zion PRA*, Chicago, 1981.
- Ellingwood, B., and T. V. Galambos, "Probability Based Criteria for Structural Design," *Structural Safety*, Vol. 1, 1982, pp. 15-26.
- Ernst, M., "NRC Plans for the Use of PRA," presented at ANS meeting, Washington, D.C., November 14-19, 1982.
- Federal Register*, 1980, Vol. 45, No. 231, Friday, November 28, 1980, pp. 79318-79378.
- Fischhoff, Baruch, Christoph Hohenemser, Roger Kasperson, and Robert Kates, "Handling Hazards," *Environment*, Vol. 20, No. 7, September 1978.
- Harriss, Robert C., and Christoph Hohenemser, "Mercury: Measuring and Managing the Risk," *Environment*, Vol. 20, No. 9, November 1978.
- Hart, G., *Uncertainty Analysis, Loads, and Safety in Structural Engineering*, Prentice Hall, New York, 1982.
- L.A. Times*, "Is The State Ready for A Big One?" Part I, p. 3, Sunday, June 27, 1982.

- L.A. Times*, "Killer Quake May Exceed Sylmar Tragedy, Scientists Say," Part II, p. 1, Sunday, July 11, 1982.
- Mantel, Nathan, and Marvin A. Schneiderman, "Estimating Safe Levels, a Hazardous Undertaking," *Career Research*, Vol. 35, June 1975, pp. 1379-1386.
- McKone, Tom, *Background Fatal Cancer Risks*, American Cancer Society, Washington, D.C., 1979.
- Meyerhoff, G. G., "Limit States Design in Geothermal Engineering," *Structural Safety*, Vol. 1, 1982, pp. 67-71.
- Morgan, M. Granger, Samuel C. Morris, Alan K. Meier, and Debra L. Shenk, "A Probabilistic Methodology for Estimating Air Pollution Health Effects from Coal-Fired Power Plants," *Energy Systems & Policy*, Vol. 2, No. 3, 1978, pp. 287-309.
- Nuclear Regulatory Commission, "Action Plan to Implement the Commission's Proposed Safety Goal Policy Statement," Draft Document, Washington, D.C., June 1982.
- Nuclear Regulatory Commission, *The Reactor Safety Study*, NUREG 75/014, WASH1400, Washington, D.C., October 1975.
- Nuclear Regulatory Commission, *Occupational Radiation Exposure at Commercial Nuclear Power Reactors*, NUREG0713, Vol. 1., Washington, D.C., January 1979.
- Nuclear Regulatory Commission, *Anticipated Transients Without Scram for Light Water Reactors*, NUREG0460, Vol. 4, Washington, D.C., March 1980.
- Nuclear Regulatory Commission, *The PRA Procedures Guide*, NUREG/CE2300, Vol. 2, Rev. 1, Washington, D.C., 1982a.
- Nuclear Regulatory Commission, "Proposed Safety Criteria for Nuclear Power Plants," NUREG0880 (Draft), Washington, D.C., June 1982b.
- Nuclear Regulatory Commission, Office of Public Affairs, News Releases for September 1982, NUREG/BR0032, ACRS letter of September 15, 1982d.
- Nuclear Regulatory Commission, "Plan To Evaluate the Commission's Safety Goal Policy Statement," Washington, D.C., January 1983.
- Okrent, D., *A Survey of Expert Opinion on Low Probability Earthquakes*, University of California at Los Angeles, UCLA-ENG-7515, February 1975.
- Philadelphia Electric Company, *Limerick Generating Station PRA*, Philadelphia, 1981.

Power Authority, State of New York, *The Indian Point PRA*, New York, 1982.

Rai, Kanta, and John Van Ryzin, "Risk Assessment of Toxic Environmental Substances Using a Generalized Multihit Dose Response Model," in N. Breslow and A. Whittemore (eds.), *Energy and Health, Proceedings of SIAM Institute*, Philadelphia, 1979, pp. 9-117.

Ricci, Paolo, "Comparative Risk Analysis of Selected Electrical Energy Systems," EPRI-RP-1772, Draft Final Report, November 1981.

Salem, Steven L., Kenneth A. Solomon, and Michael S. Yesley, *Issues and Problems in Inferring a Level of Acceptable Risk*, The Rand Corporation, R-2561-DOE, August 1980.

Scott, Stanley, *Policies for Seismic Safety: Elements of a State Governmental Program*, Institute of Governmental Studies, University of California at Berkeley, 1979.

Solomon, K. A., and S. C. Abraham, "A Useful Measure for Comparing Occupational Risk Across Industries," *Health Physics*, Vol. 38, pp. 375-391, 1980.

Solomon, K. A., M. Meyer, and P. F. Nelson, *Management of Risks Associated with Drinking Water at the State and Local Levels*, University of California at Los Angeles, UCLA School of Engineering and Applied Science, UCLA Eng.-8243, May 1982.

Solomon, Kenneth A., and David Okrent, *Seismic Building Codes for the City of Los Angeles, California--A Brief Case Study*, The Rand Corporation, P-6018, November 1977.

The Engineer of California, August 1977, Vol. 30, No. 8, August 1977, p. 9.

U.S. Senate, *Food Safety: Where Are We?* Prepared for the Committee on Agriculture, Nutrition and Forestry, Washington, D.C., 1979.

Whittemore, Alice, "Quantitative Methods," *Environmental Health Siam* Philadelphia, 1977.

